Parametrization of Generalized Primal-Dual Triangulations

Pooran Memari^{1,2}, Patrick Mullen¹, and Mathieu Desbrun¹

- ¹ California Institute of Technology
- ² CNRS LTCI, Télécom ParisTech

Abstract

Motivated by practical numerical issues in a number of modeling and simulation problems, we introduce the notion of a *compatible dual complex* to a primal triangulation, such that a simplicial mesh and its compatible dual complex (made out of convex cells) form what we call a *primal-dual triangulation*. Using algebraic and computational geometry results, we show that compatible dual complexes exist only for a particular type of triangulation known as weakly regular. We also demonstrate that the entire space of primal-dual triangulations, which extends the well known (weighted) Delaunay/Voronoi duality, has a convenient, geometric parametrization. We finally discuss how this parametrization may play an important role in discrete optimization problems such as optimal mesh generation, as it allows us to easily explore the space of primal-dual structures along with some important subspaces.

1 Introduction

Mesh generation traditionally aims at tiling a bounded spatial domain with simplices (triangles in 2D, tetrahedra in 3D) so that any two of these simplices are either disjoint or sharing a lower dimensional face. The resulting triangulation provides a discretization of space through both its primal (simplicial) elements and its dual (cell) elements. Both types of element are crucial to a variety of numerical techniques, finite element (FE) and finite volume (FV) methods being arguably the most widely used in computational science. To ensure numerical accuracy and efficiency, specific requirements on the size and shape of the primal (typically for FE) or the dual elements (typically for FV) in the mesh are often sought after.

Primal/Dual Pairs. A growing trend in numerical simulation is the simultaneous use of primal and dual meshes: Petrov-Galerkin finite-element/finite-

volume methods (FE/FVM, [BP83, McC89, PMH07]) and exterior calculus based methods [Bos98, DKT07, GP10] use the ability to store quantities on both primal and dual elements to enforce (co)homological relationships in, e.g., Hodge theory. The choice of the dual, defined by the location of the dual vertices, is however not specified a priori. A very common dual to a triangulation in \mathbb{R}^d is the cell complex which uses the circumcenters of each d-simplex as dual vertices. If the initial triangulation is Delaunay (i.e., satisfying the empty circumsphere property [Ede87]), this dual is simply the Voronoi diagram of the primal vertices, and its nice properties of non-self-intersection, convexity, and orthogonality of the primal and dual elements have led to its use in countless papers in computational sciences. The barycentric dual, for which barycenters are used instead of circumcenters, is used for certain finite-volume computations, but it fails to satisfy both the orthogonality and (more importantly) convexity properties on general triangulations.

Towards Generalized Primal/Dual Meshes. While the Delaunay-Voronoi duality [PS85, Ede87] is one of the cornerstones of meshing methods and, as such, has been extensively used in diverse fields, more general dualities are often desired. In biology for example, Voronoi cells (along with their dual triangulations) were initially identified as closely approximating a variety of monolayer cells and epithelial cells in tissue [OBSC00]; but computational biologists are now seeking generalizations of Voronoi diagrams to parameterize a larger set of convex polyhedral tilings [Mjo06] to model the development of early animal tissues and shoot meristems. The Weighted-Delaunay/Laguerre duality was also advocated recently in [MMdGD11] to help provide lower error bounds on computations. Building on Schlegel diagrams and a number of results in algebraic and computational geometry, we present an even more general primal-dual pairs of complexes that we denote as primal-dual triangulations.

Contributions. While Delaunay/Voronoi or Weighted-Delaunay/Laguerre duality assumes orthogonality between primal and dual complexes, we relax this requirement in our work: we investigate a general notion of dual complex to a triangulation which we call a compatible dual complex, where the dual complex is only assumed to be a union of convex straight-edge polytopes, with an adjacency graph matching the triangulation's adjacency graph—but where k-simplices of the triangulation are not necessarily orthogonal to their associated (n-k)-cells of the dual complex. While we will show that any two-dimensional triangulation admits a compatible dual complex, this property is no longer true in dimension three and above. We introduce a proper characterization of primal-dual triangulations, which results in a simple parametrization of the whole space of primal-dual triangulations. Finally, we discuss potential applications of our contributions, for instance in mesh optimization and clustering, as our parametrization allows us to easily explore a space of primal-dual structures much larger than the space of orthogonal primal-dual structures.

2 Preliminaries

Before delving into our contributions, we start by reviewing important notions related to triangulations in different fields (such as combinatorial, computational, and algebraic geometry) that we build upon and extend in subsequent sections.

2.1 Complex, Subdivision, and Triangulation

A **cell complex** in \mathbb{R}^d is a set K of convex polyhedra (called cells) satisfying two conditions:

- 1. Every face of a cell in K is also a cell in K, and
- 2. If C and C' are cells in K, their intersection is either empty or a common face of both.

A simplicial complex is a cell complex whose cells are all simplices. The **body** |K| of a complex K is the union of all its cells. When a subset P of \mathbb{R}^d is the body of a complex K, then K is said to be a **subdivision** of P; if, in addition, K is a simplicial complex, then K is said to be a **triangulation** of P. For a set X of points in \mathbb{R}^d , a triangulation of X is a simplicial complex K for which each vertex of K is in X.

Note that in the definition of a triangulation of X, we do not require all the points of X to be used as vertices; a point $\mathbf{x_i} \in X$ is called *hidden* if it is not used in the triangulation. A triangulation of X with no hidden points is called a *full* triangulation of X.

2.2 Triangulations in \mathbb{R}^d through Lifting in \mathbb{R}^{d+1}

Let $\mathbf{X} = \{\mathbf{x_1}, \dots, \mathbf{x_n}\}$ be a set of points in \mathbb{R}^d . A simple way of constructing a triangulation of \mathbf{X} is through the following *lifting procedure*: take an arbitrary function $L: \mathbf{X} \longrightarrow \mathbb{R}$ called the *lifting function*; consider the points $(\mathbf{x_i}, L(\mathbf{x_i})) \in \mathbb{R}^{d+1}$, i.e., the points of \mathbf{X} *lifted* onto the graph of L; in the space \mathbb{R}^{d+1} , consider $\operatorname{Conv}(L)$ the convex hull of vertical rays $\{(x_i, l) | l \geq L(\mathbf{x_i}), l \in \mathbb{R}, x_i \in X\}$; the bounded faces of $\operatorname{Conv}(L)$, i.e. faces which do not contain vertical half lines, form the **lower envelope** of the lifting L. If the function L is generic (see [GKZ94] Chap. 7), the orthogonal projection (onto the first d coordinates) of the lower envelope of L produces a triangulation of \mathbf{X} .

It is clear that the above lifting procedure may produce triangulations for which not all points of \mathbf{X} are vertices. A triangulation of a set \mathbf{X} of points obtained through lifting is full (i.e., has no hidden points) if and only if all the points $(\mathbf{x_i}, L(\mathbf{x_i}))$ lie on the **lower envelope of** L (or, in other words, if function L can be extended, through linear interpolation in the triangles, to a **convex** piecewise-linear function).

2.3 Regular Triangulations and Subdivisions

Let **X** be a finite set of points in \mathbb{R}^d . A triangulation obtained by orthogonally projecting the lower envelope of a lifting of **X** in \mathbb{R}^{d+1} onto the first d coordinates is called a regular triangulation ([Zie95], Definition 5.3).

More generally, a subdivision T of a polytope $P \subset \mathbb{R}^d$ is regular if it arises from a polytope $Q \subset \mathbb{R}^{d+1}$ in the following way:

- The polytope P is the image $\pi(Q) = P$ of the polytope Q via the projection that deletes the last coordinate.
- The complex T is the projection under π of all the lower faces of Q. We call F a lower face of Q if for every $\mathbf{x} \in F$ and a real number $\lambda > 0$, $\mathbf{x} \lambda \mathbf{e}_{d+1} \notin Q$. Informally, they are the faces that you can see from P if you "look up" at Q.

Regular triangulations have appeared in different mathematical contexts and are known under different names, such as weighted Delaunay triangulations (see next section) or coherent triangulations [GKZ94]. While every point set admits regular triangulations (see [DLRS], Proposition 2.2.4.), not every triangulation is regular; Figure 4 (left) illustrates a classical non-regular triangulation.

2.4 Schlegel diagrams

Schlegel diagrams are a related mathematical notion based on a construction very similar to the lifting procedure defined above. They have been proven an important tool for studying combinatorial and topological properties of polytopes, as well as for visualizing four-dimensional polytopes [Ban90].

A Schlegel diagram is the (perspective) projection of a polytope from \mathbb{R}^{d+1} to \mathbb{R}^d through a point p outside of the polytope, above the center of a facet f. All vertices and edges of the polytope are projected onto the hyperplane of that facet in the following way: For any vertex of the polytope, the line from p through the vertex meets the hyperplane of f at the image of the vertex. If two vertices are connected by an edge in the polytope, then the image of the edge is the segment joining the images of the two vertices.

If the polytope is convex, there exists a point near the facet f which maps all other facets inside f, so no edges need to cross in the projection. In other words, the Schlegel diagram of a convex polytope is the *projection* of the polytope's skeleton on one of its faces (the nodes corresponding to the vertices which do not belong to that face must lie inside the face). In this case the resulting entity is a polytopal subdivision of the facet in \mathbb{R}^d that is combinatorially equivalent to the original polytope, see Figure 1.

Intuitively, by sending the base point of a Schlegel diagram at infinity we obtain a regular subdivision (see [Zie95] Prop. 5.9, or [Tho06]). We will use this equivalence between regular subdivisions and Schlegel diagrams (illustrated in Figure 2) later in this paper.

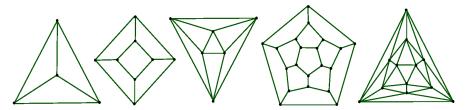


Fig. 1: The Schlegel diagrams of the five polyhedra (tetrahedron, cube, octahedron, dodecahedron and icosahedron).



Fig. 2: By sending the base point of a Schlegel diagram at infinity we obtain a regular subdivision (after [Sch03]).

Proposition 1 ([Zie95]). If T is a Schlegel diagram, T is a regular subdivision of |T|.

2.5 Weighted Delaunay Triangulations

A special choice for the lifting function produces the well-known and widely-used Delaunay triangulation (see [Raj94, Mus97] for their properties, and [PS85] for numerous applications). Indeed, let X be a set of points in \mathbb{R}^d . Consider the lifting of the points in X onto the surface of the paraboloid $h(\mathbf{x}) = \|\mathbf{x}\|^2$ in \mathbb{R}^{d+1} ; i.e., each $\mathbf{x_i} = (a_1, \dots, a_d) \in X$ gets mapped to $(\mathbf{x_i}, h_i) \in \mathbb{R}^{d+1}$ with $h_i = \|\mathbf{x_i}\|^2 = a_1^2 + \dots + a_d^2$. Then the orthogonal projection of the lower envelope of this lifting produces a (full) triangulation coinciding with the **Delaunay triangulation** of X.

A regular triangulation can now be seen as a generalization of the Delaunay triangulation as follows. We first define a **weighted point set** as a set $(\mathbf{X}, W) = (\mathbf{x_1}, w_1), \dots, (\mathbf{x_n}, w_n)$, where **X** is a set of points in \mathbb{R}^d , and $\{w_i\}_{i \in [1, \dots, n]}$ are real numbers called weights. The **weighted Delaunay triangulation** of (\mathbf{X}, W) is then the triangulation of **X** obtained by projecting the lower envelope of the points $(\mathbf{x_i}, \|\mathbf{x_i}\|^2 - w_i) \in \mathbb{R}^{d+1}$. Note that a weighted Delaunay triangulation can now have hidden points.

Notice also that given a lifting function L and its values $l_i = L(\mathbf{x_i})$ at the points of X, one can always define weights to be the difference between the paraboloid and the function L, $w_i = ||\mathbf{x_i}||^2 - l_i$. We conclude that the notions of regular triangulations and weighted Delaunay triangulations are *equivalent*.

2.6 Generalized Voronoi Diagrams (Power Diagrams)

Delaunay triangulations (resp., weighted Delaunay triangulations) can also be obtained (or defined) from their dual *Voronoi diagrams* (resp., power diagrams). Let $(\mathbf{X}, W) = \{(\mathbf{x_i}, w_i)\}_{i \in I}$ be a weighted point set in \mathbb{R}^d . The power of a point $\mathbf{x} \in \mathbb{R}^d$ with respect to a weighted point $(\mathbf{x_i}, w_i)$ (sometimes referred to as the Laguerre distance) is defined as $d^2(\mathbf{x}, \mathbf{x_i}) - w_i$, where d(.,.) stands for the Euclidean distance. Using this power definition, to each x_i we associate its weighted Voronoi region $V(\mathbf{x_i}, w_i) = \{\mathbf{x} \in \mathbb{R}^d | d^2(\mathbf{x}, \mathbf{x_i}) - w_i \leq d^2(\mathbf{x}, \mathbf{x_j}) - w_j, \forall j\}$. The power diagram of (\mathbf{X}, W) is the cell complex whose cells are the weighted Voronoi regions.

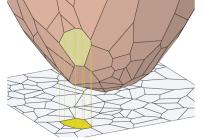
Note that when the weights are all equal, the power diagram coincides with the Euclidean Voronoi diagram of X. Power diagrams are well known to be dual to weighted Delaunay triangulations, as we review next.

2.7 Power Diagram vs. Weighted Delaunay Triangulation

The **dual** of the power diagram of (\mathbf{X}, W) is the weighted Delaunay triangulation of (\mathbf{X}, W) . This triangulation contains a k-simplex with vertices $\mathbf{x_{a_0}}, \mathbf{x_{a_1}}, \dots, \mathbf{x_{a_k}}$ in \mathbf{X} if and only if $V(x_{a_0}, w_{a_0}) \cap V(x_{a_1}, w_{a_1}) \cap \cdots \cap V(x_{a_k}, w_{a_k}) \neq \emptyset$, $\forall k \geq 0$. While many other generalization of Voronoi diagrams exist, they do not form straight-edge and convex polytopes, and are thus not relevant here.

Geometric Property. We finally review an interesting geometric property of Voronoi and power diagrams, as it will become useful later on. Consider the affine functions $\sigma_i(x) = -2\mathbf{x_i} \cdot \mathbf{x} + ||\mathbf{x_i}||^2$ for i = 1, ..., n: their graphs are obviously hyperplanes of \mathbb{R}^{d+1} that are tangent to the paraboloid

 $h(\mathbf{x}) = \|\mathbf{x}\|^2$ at point $\mathbf{x_i}$. Let us call these hyperplanes H_i , and let H_i^+ denote the half-space lying above H_i . The minimization diagram of the σ_i is obtained by projecting the lower envelop of $H_1^+ \cap \cdots \cap H_n^+$ orthogonally onto \mathbb{R}^d . However, for any \mathbf{x} , $\operatorname{argmin}_i \|\mathbf{x} - \mathbf{x_i}\|^2 = \operatorname{argmin}_i (-2\mathbf{x_i} \cdot \mathbf{x} + \|\mathbf{x_i}\|^2)$. Thus, one concludes that the Euclidean Voronoi dia-



gram of \mathbf{X} is the orthogonal projection of the skeleton of the lower envelop of $H_1^+ \cap \cdots \cap H_n^+$ —see inset and [ES86]. More generally, the power diagram of (\mathbf{X}, W) is the orthogonal projection of the skeleton of the lower envelop of the intersection of half-spaces lying above hyperplanes $(-2\mathbf{x_i} \cdot \mathbf{x} + \|\mathbf{x_i}\|^2 - w_i)$: such a hyperplane is indeed the tangent plane to the paraboloid $h(\mathbf{x}) = \|\mathbf{x}\|^2$ at point $\mathbf{x_i}$ translated down by w_i . This geometric fact can be interpreted as a hyperplane assignment to each (weighted) Voronoi region, and it will be used in the forthcoming proof of our Theorem 1.

3 Compatible Dual Complexes of Triangulations

We now show that the notion of mesh duality can be extended so that the dual complex is defined geometrically, and independently from the triangulation—while the combinatorical compatibility between the triangulation and its dual is maintained.

Definition 1 (Simple Cell Complex). A cell complex K in \mathbb{R}^d is called **simple** if every vertex of K is incident to d+1 edges. K is called labeled if every d-dimensional cell of K is assigned a unique label; in this case, we write $K = \{C_1, \ldots, C_n\}$, where n is the number of d-dimensional cells of K, and C_i is the i-th d-dimensional cell.

Definition 2 (Compatible Dual Complex). Let T be a triangulation of a set $\mathbf{X} = \{\mathbf{x_1}, \dots, \mathbf{x_n}\}$ of points in \mathbb{R}^d , and $K = \{C_{i_1}, \dots, C_{i_n}\}$ be a labeled simple cell complex, i.e. there is a one-to-one correspondence between x_p and C_{i_p} . K is called a **compatible dual complex** of T if, for every pair of points $\mathbf{x_p}$ and $\mathbf{x_q}$ that are connected in T, C_{i_p} and C_{i_q} share a face.

This compatibility between K and T is purely combinatorial, i.e., it simply states that the connectivity between points induced by K coincides with the one induced by T. Notice that the cell C_{i_p} associated to the point $\mathbf{x_p}$, does not necessarily contain $\mathbf{x_p}$ in its interior. Moreover, the edge $[\mathbf{x_p}, \mathbf{x_q}]$ and its $dual\ C_{i_p} \cap C_{i_q}$ are not necessarily orthogonal to each other, unlike most conventional geometric dual structures. Consequently, we can generalize the notion of mesh duality through the following definition:

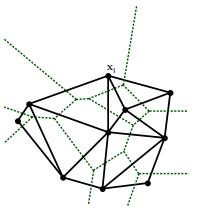
Definition 3 (Primal-Dual Triangulation (PDT)). A pair (T, K) is said to form a d-dimensional **primal-dual triangulation** if T is a triangulation in \mathbb{R}^d and K is a compatible dual complex of T. If every edge $[\mathbf{x_p}, \mathbf{x_q}]$ and its dual $C_{i_p} \cap C_{i_q}$ are orthogonal to each other, the pair (T, K) is said to form an **orthogonal primal-dual triangulation**.

3.1 Characterization of Primal-Dual Triangulations

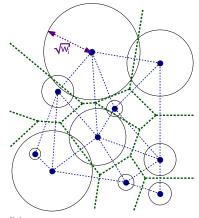
An immediate question is whether any triangulation can be part of a PDT. We first characterize the triangulations that admit a compatible dual complex through the following two definitions:

Definition 4 (Combinatorial Equivalence). Two triangulations T and T' are combinatorially equivalent if there exists a labeling which associates to each point $\mathbf{x_i}$ in T a point $\mathbf{x_i'}$ in T' so that the connectivity between $\mathbf{x_i}$'s induced by T matches the connectivity between the $\mathbf{x_i'}$'s induced by T'.

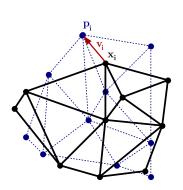
Definition 5 (Combinatorially Regular Triangulations (CRT)). A triangulation T of a d-dimensional point set X is called a **combinatorially regular triangulation** if there exists a d-dimensional point set X' admitting a regular triangulation T' such that T and T' are combinatorially equivalent.



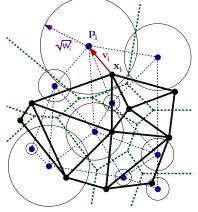
(a) A triangulation T (solid line) and a compatible (non-orthogonal) dual complex K (dashed line).



(b) K is the power diagram dual to a regular triangulation T' (thin dashed line).



(c) The triangulation T is combinatorially equivalent to the regular triangulation T'.



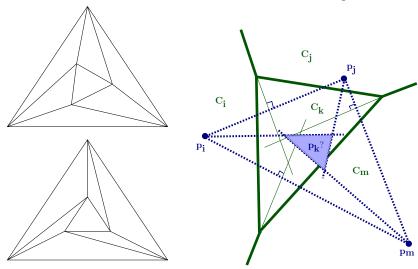
(d) PDT (T, K) is parameterized through $\mathbf{x_i}$ (points), $\mathbf{v_i}$ (displacements), and w_i (weights).

Fig. 3: Primal-Dual Triangulation, with its primal triangulation, dual complex, and combinatorially equivalent regular triangulation separately displayed for clarity.

Remark: these CRT triangulations have been introduced in [Lee91] under the name of weakly regular triangulations, since a displacement of their vertices suffices to make them regular. Figure 4 (after [Lee91]) shows an example of a combinatorially regular triangulation which is not, itself, regular.

Existence of PDTs in 2D. The 2D case is rather simple, due to this result:

Proposition 2 ([Lee91]). Any 2-dimensional triangulation is combinatorially regular.



- (a) A regular triangulation (top), once deformed (bottom), becomes a combinatorially regular triangulation which is *not*, itself, regular.
- (b) A cell complex (solid line) that does not admit a primal orthogonal triangulation (after [Aur87a]).

Fig. 4: Classical Examples, showing the existence of non-regular triangulations (a) and of dual complexes without associated primal triangulations (b).

Proof. This is a straightforward corollary of a classical theorem of Steinitz [Ste22] which implies that for every complex in the plane whose edge graph is three-connected, there exists a convex 3-polyhedron with isomorphic boundary. Since any 2-dimensional triangulation T is trivially edge three-connected, there exists a convex 3-polyhedron P with isomorphic boundary. Therefore, the orthogonal projection of the boundary of P onto the plane is a regular triangulation which is isomorphic to T. \Box

Therefore, every 2D triangulation T can be part of a PDT pair (T, K).

Existence of PDTs in Higher Dimensions. In higher dimensions (three and above), however, the situation is rather different: the set of regular triangulations does not contain all possible triangulations between points as stated in the next proposition.

Proposition 3 ([Grü03]). For $d \geq 3$, there exist d-D triangulations which are not combinatorially regular.

This is equivalent to the fact that there are simplicial diagrams that are not combinatorially equivalent to any Schlegel diagram. Chronologically, the first example was found by Grünbaum in 1965 (see [Grü03] pages 219-224): he presented a simplicial 3-diagram with 7 vertices which is not combinatorially equivalent to any Schlegel diagram of 4-polytopes. Before this result, there

had been several attempts to prove that any simplicial diagram is combinatorially equivalent to a Schlegel diagram (also called polytopal or polytopical)³, see also [CH80]. The simplest non-combinatorially regular examples are the Brucker sphere and the Barnette sphere (short proofs are available in [Ewa96]Sec III.4).

3.2 PDT=CRT

We now show that combinatorially regular triangulations are the only ones that admit compatible dual complexes. The proof revolves around a theorem due to Aurenhammer:

Every simple cell complex in \mathbb{R}^d , $d \geq 3$, is a Schlegel diagram.⁴

This theorem was proved in [Aur87a] through an iterative construction which is valid in any dimension $d \geq 3$. We use this theorem to prove the following theorem which surprisingly implies that in higher dimensions there are triangulations that do *not* admit a dual complex:

Theorem 1 (PDT Characterization). A d-dimensional triangulation T admits a compatible (not necessarily orthogonal) dual complex if and only if T is combinatorially regular.

Proof. Let T be a triangulation combinatorially equivalent to a regular triangulation T'. The conventional dual of T', a power diagram, is thus a compatible dual of T as well.

Now suppose that T is a triangulation with a compatible dual complex K. According to Aurenhammer's theorem, K is a Schlegel diagram, i.e., to each cell C_j of K we can assign a hyperplane H_j in \mathbb{R}^{d+1} so that the lower envelop of these hyperplanes projects orthogonally onto K (see also [Sha93]). On the other hand, the tangent planes of the (d+1)-dimensional paraboloid span the whole space of non vertical hyperplane directions in \mathbb{R}^{d+1} . Therefore for each j, there exists a point $\mathbf{p_j} \in \mathbb{R}^d$ whose tangent plane to the paraboloid in \mathbb{R}^{d+1} is parallel to H_j . In that case, one can easily conclude that K is the power diagram of $(\mathbf{p_j}, w_j), 1 \leq j \leq n$, where the weight w_j depends on the distance between the two parallel hyperplanes mentioned before. Now consider the weighted Delaunay triangulation T' of weighted points $(\mathbf{p_j}, w_j), 1 \leq j \leq n$. Since K is a compatible dual complex for both T and T', we conclude that T and T' are combinatorially equivalent. Since T' is regular, T is a combinatorially regular triangulation. \square

³These terms are commonly used to study the combinatorics of polytopes and diagrams, often independently from their embedding. Therefore, the reader should be aware of some ambiguities in the definition of these notions in the literature, between being a Schlegel diagram or being combinatorially equivalent to it.

⁴Note that Aurenhammer employs the equivalent name of polytopical diagram.

Existence of PDT for Given Dual Complexes. We can also discuss whether a dual complex can be part of a PDT pair to further characterize the space of primal-dual pairs we will be parameterizing. In dimension $d \geq 3$, Aurenhammer's theorem directly implies that every simple cell complex is dual to a regular triangulation. But this same theorem is not true in 2D, and Figure 4 (right) shows a counter example presented by Aurenhammer [Aur87a]—proving that not every 2D simple cell complex is the dual of a regular triangulation; thus PDTs do not capture all 2D convex cell complexes.

4 Parameterizing Primal-Dual Triangulations

We have established that primal-dual triangulations cover all dual complexes in $d \geq 3$ but only those which come from a regular triangulation for d = 2; they also cover all 2D triangulations, but only triangulations which admit a dual in $d \geq 3$. We now focus on parameterizing the whole space of primal-dual triangulations with n points in \mathbb{R}^d by simply adding parameters at the points. We then explore a geometric interpretation of this intrinsic parametrization as well as its properties.

The proof of Theorem 1 leads us very naturally to a parametrization of all the triangulations that admit a compatible dual complex:

Definition 6. A parameterized primal-dual triangulation is a primal-dual triangulation parameterized by a set of triplets $(\mathbf{x_i}, w_i, \mathbf{v_i})$, where $\mathbf{x_i}$ is the position in \mathbb{R}^d of the i^{th} node, w_i is a real number called the weight of x_i , and $\mathbf{v_i}$ is a d-dimensional vector called the displacement vector of $\mathbf{x_i}$. The triangulation associated with the triplets $(\mathbf{x_i}, w_i, \mathbf{v_i})$ is defined such that its dual complex K is the power diagram of weighted points $(\mathbf{p_i}, w_i)$, where $\mathbf{p_i} = \mathbf{x_i} + \mathbf{v_i}$.

The dual complex K can be seen as the generalized Voronoi diagram of the $\mathbf{x_i}$'s for the distance $d(\mathbf{x}, \mathbf{x_i}) = ||\mathbf{x} - \mathbf{x_i} - \mathbf{v_i}||^2 - w_i$. When the vectors $\mathbf{v_i}$ are all null, the parameterized primal-dual triangulation T is regular, thus perpendicular to its dual K, and the pair (T, K) forms an orthogonal primal-dual triangulation. This proves that weighted Delaunay triangulations are sufficient to parameterize the set of all orthogonal primal-dual triangulations (see also [Gli05]). The displacement vectors extend the type of triangulations and duals we can parameterize.

Geometric Interpretation. This parametrization can be seen as a very natural extension of the geometric interpretation mentioned in Section 2.6. Indeed, PDTs in \mathbb{R}^d are simply defined by assigning one hyperplane H_i in \mathbb{R}^{d+1} per point x_i : the orthogonal projection of the lower envelop of these hyperplanes $\{H_i\}_{i=1,...,n}$ will form the dual complex of the PDT, inducing the primal triangulation. Weights and displacements vectors serve as a means to encode the choice of hyperplanes.

We now provide some geometric properties of the primal-dual triangulation parameterized by a set of triplets $(\mathbf{x_i}, w_i, \mathbf{v_i})$.

Lemma 1 (Characterization of PDT triplets). Let (T, K) be a PDT parameterized by $(\mathbf{x_i}, w_i, \mathbf{v_i})$. The set of triplets $\{(\mathbf{x_i'}, w_i', \mathbf{v_i'})\}$ parameterizes the *same* primal-dual triangulation (T, K) if and only if $\mathbf{x_i'} = \mathbf{x_i}$ and there exist a constant $\alpha \in \mathbb{R}$, a positive constant $\beta \in \mathbb{R}^+$, and a vector \mathbf{v} such that $w_i' = w_i - 2\mathbf{v} \cdot (\mathbf{x_i} + \mathbf{v_i}) + \alpha$ and $\mathbf{v_i'} = \beta(\mathbf{x_i} + \mathbf{v_i}) - \mathbf{x_i} + \mathbf{v}$.

Proof. Let us denote by R and R' the weighted Delaunay triangulations of weighted points $(\mathbf{x_i} + \mathbf{v_i}, w_i)$ and $(\mathbf{x_i} + \mathbf{v_i'}, w_i')$ respectively. K is then the power diagram of both weighted point sets $(\mathbf{x_i} + \mathbf{v_i}, w_i)$ and $(\mathbf{x_i} + \mathbf{v_i'}, w_i')$, and is dual to both R and R'. However, two regular triangulations have the same power diagram only if they are homothetic (see, for instance, [Aur87b] for a proof). Hence, there exist a positive constant $\beta > 0$, and a vector \mathbf{v} such that $\mathbf{x_i} + \mathbf{v_i'} = \beta(\mathbf{x_i} + \mathbf{v_i}) + \mathbf{v}$. On the other hand, it is easy to show that the power diagram of the set of weighted points $(\mathbf{p_i}, w_i)$, translated by a vector $-\mathbf{v}$, coincides with the power diagram of $(\mathbf{p_i}, w_i - 2\mathbf{v} \cdot \mathbf{p_i} + \alpha)$. Also, such a translated power diagram can be changed back to the original diagram by simply adding \mathbf{v} to each displacement vector $\mathbf{v_i}$. This implies that $(\mathbf{x_i}, w_i - 2\mathbf{v} \cdot (\mathbf{x_i} + \mathbf{v_i}) + \alpha$, $(\beta - 1)\mathbf{x_i} + \beta\mathbf{v_i} + \mathbf{v})$ and $(\mathbf{x_i}, w_i, \mathbf{v_i})$ parameterize the same PDT. \square

This lemma characterizes the classes of *equivalent* triplets parameterizing the same PDT. Using this characterization, we now provide a set of constraints for the parameters of the triplets which allow us to avoid redundancy between equivalent triplets and define an efficient parametrization for the space of primal-dual triangulations:

Theorem 2 (PDT Parametrization). There is a bijection between all primal-dual triangulations in \mathbb{R}^d and sets of triplets $(\mathbf{x_i}, w_i, \mathbf{v_i})$, $1 \leq i \leq n$, where $\mathbf{x_i}, \mathbf{v_i} \in \mathbb{R}^d$, $w_i \in \mathbb{R}$ with $\sum_i w_i = 0$, $\sum_i \mathbf{v_i} = 0$, and $\sum_i \|\mathbf{x_i} + \mathbf{v_i}\|^2 = \sum_i \|\mathbf{x_i}\|^2$.

Proof. The proof is provided in two parts: i) For any triplet $(\mathbf{x_i}, w_i, \mathbf{v_i})$, there exists an equivalent triplet $(\mathbf{x_i'}, w_i', \mathbf{v_i'})$ which fulfills the conditions of the theorem. Indeed, the choice of:

$$\beta = \sqrt{\frac{n\sum_{i} \|\mathbf{x_i}\|^2 - \|\sum_{i} \mathbf{x_i}\|^2}{n\sum_{i} \|\mathbf{x_i} + \mathbf{v_i}\|^2 - \|\sum_{i} (\mathbf{x_i} + \mathbf{v_i})\|^2}},$$
$$\mathbf{v} = \frac{1 - \beta}{n} \sum_{i} \mathbf{x_i} - \frac{\beta}{n} \sum_{i} \mathbf{v_i}, \quad \alpha = \frac{2}{n} \mathbf{v} \cdot \sum_{i} (\mathbf{x_i} + \mathbf{v_i}) - \frac{\sum_{i} w_i}{n}$$

in the characterization of Lemma 1 gives the desired triplet. ii) Suppose that both (x_i, w_i, v_i) and $(\mathbf{x_i}, w_i - 2\mathbf{v} \cdot \mathbf{x_i} - 2\mathbf{v} \cdot \mathbf{v_i} + \alpha, (\beta - 1)\mathbf{x_i} + \beta \mathbf{v_i} + \mathbf{v})$ fulfill the conditions, for some constants α and $\beta > 0$, and a constant vector \mathbf{v} . A

direct computation implies $\alpha = 0$, $\mathbf{v} = 0$ and $\beta = 1$. Therefore for each class of equivalent triplets, there is a single triplet which verifies the conditions of the parametrization. \square

Remark: using this parametrization, the particular case of Delaunay / Voronoi PDT of a set of points $\{\mathbf{x_i}\}_{i=1..n}$ is naturally parameterized by triplets $(\mathbf{x_i}, 0, 0)$. Similarly, the weighted Delaunay / Power PDT is parameterized by triplets $(\mathbf{x_i}, w_i, 0)$. Note also that the condition $\sum_i w_i = 0$ may be replaced by $\min_i w_i = 0$, by simply subtracting the minimum weight from all the weights of triplets. This new condition implies that all the weights are positive, which may be useful in some applications.

5 Conclusions

In this paper, we introduced the notion of compatible dual complex for a given triangulation in \mathbb{R}^d , and discussed the conditions under which an arbitrary triangulation admits a compatible, possibly non-orthogonal dual complex (and vice-versa). Note that our only assumption on the dual is that it is made out of *convex polytopes*, thus reducing the space of possible primal-dual pairs to a computationally-convenient subset for which basis functions and positive barycentric coordinates are easily defined. We also pointed out a link to a previously-introduced notion of weakly regular triangulation by Lee in the nineties, and that there are triangulations that do *not* admit a dual complex. We derived a natural parametrization of all non-orthogonal primal-dual structures in \mathbb{R}^d by means of plane assignments in \mathbb{R}^{d+1} .

Besides the theoretical interest of these new primal-dual structures, we anticipate numerous applications. We believe that our results can benefit mesh optimization algorithms as we provide a particularly convenient way to explore a large space of primal-dual structures. We recently provided a first step in this direction by designing pairs of primal-dual structures that optimize accuracy bounds on differential operators using our parametrization [MMdGD11], thus extending variational approaches designed to improve either primal (Optimal Delaunay Triangulations [ACSYD]) or dual (Centroidal Voronoi Tessellations [LWL⁺09]) structures. To some extent, our approach can even help to deal with situations where the primal triangulation is given and cannot safely be altered: for instance, moving vertices and/or changing the connectivity of a triangle mesh in \mathbb{R}^3 is potentially harmful, as it affects the surface shape. Still, the ability to optimize weights to drive the selection of the dual mesh is very useful. We can easily optimize primal-dual triangulations (meshes) by minimizing a functional (energy) with respect to weights. The connectivity is kept intact, regardless of the weights—only the position and shape of the compatible dual is optimized. Our 2D and 3D experiments [MMdGD11] show that only optimizing the weights is particularly simple and beneficial on a number of meshes. Fig. 5 depicts a triangle mesh of a hand and its intrinsic dual before and after weight optimization, showing a drastic reduction in the number of negative dual edges—thus providing a practical alternative to the use of intrinsic Delaunay meshes advocated in [FSBS06].

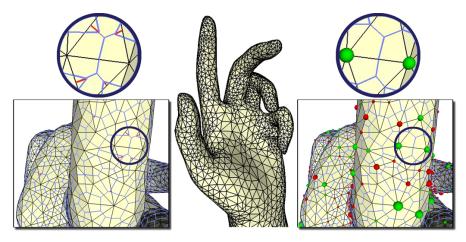


Fig. 5: Improving Dual Structure of a Surface Mesh: For a given triangulation (center) there are several triangles whose circumcenter is far outside the triangle (left, lines in red). By optimizing only the weights the new dual vertices are better placed inside the unchanged triangles (right) while keeping primal/dual orthogonality.

As another illustrative example, Fig. 6 shows that even an optimized Delaunay triangulation (ODT mesh) with exceptionally high-quality tetrahedra [TWAD09] can be made significantly better centered (i.e., with dual vertices closer to the inside of their associated primal simplex) using a simple weight optimization. Note also that in this example the number of tetrahedra with a dual vertex outside of the primal tet dropped from 17041 on the ODT mesh to 5489 on the optimized mesh—a two third reduction of outcentered tetrahedra. While all these results only explored the orthogonal primal-dual triangulations, we expect that better results would be obtained with our parametrization if we relax the orthogonality constraint in favor of arbitrary convex dual cells. The use of vertex weights to ensure boundary and feature protection as proposed in [CDL08] would be interesting as well in this context. In addition to applications in mesh optimization, as we mentioned in the introduction, modeling (as in computational biology) that uses convex space tilings could directly use our parametrization of PDTs. Clustering techniques based on k-means may also benefit from parameterizing clusters by more than just centers, as weights and vectors add more flexibility to the segmentation of input data.

Future Work. Enforcing proper embedding of a PDT can be crucial in some applications. While sufficient conditions on the weights $\{w_i\}$ and displacements $\{\mathbf{v_i}\}$ are easy to derive (one can, for instance, limit each $\mathbf{v_i}$ to stay within the ball centered on $\mathbf{x_i}$ and of radius $\sqrt{w_i}$), it could be beneficial

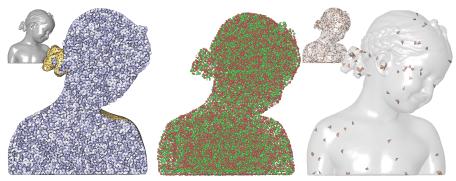


Fig. 6: Improving Dual Structures of 3D Meshes: The dual of a high-quality ODT mesh of the Bimba con Nastrino (left, cross-section; 195K tets, 36K vertices) can be optimized by improving minimal dual edge length and self-centeredness [MMdGD11] (middle; weights are displayed according to sign (red/green) and magnitude (radius)). When we single out the tetrahedra with a distance between weighted circumcenter and barycenter greater than 0.5% of the bounding box, one can see the optimized mesh (right) is significantly better than the original ODT (right, inset), even if the primal triangulations are exactly matching.

to have less constraining conditions. Other necessary and sufficient conditions to enforce, for instance, that primal vertices are placed within their associated dual cells (or vice-versa) could be also useful. Finally, we wish to study possible links between primal-dual triangulations and an algebraic-geometric construction due to Gelfand, Kapranov, and Zelevinsky [GKZ94], where each (embedded) triangulation of a set of n points $\mathbf{x_1}, \dots, \mathbf{x_n}$ in \mathbb{R}^d is associated with a point in \mathbb{R}^n . The *i*-th coordinate of this new point is the total volume of all simplices incident to x_i in the triangulation, and the convex hull of these new points is called the secondary polytope of the point set. This hull has dimension n-d-1, and interestingly, its vertices correspond exactly to the regular triangulations of the point set. This construction helps translate geometric questions into combinatorial questions about polyhedral fans. Using a similar construction for combinatorially regular triangulations could extend some existing results for regular triangulations—in particular, algorithmic and enumerative questions. For instance, the decision of whether a given triangulation is regular is easily reduced to linear programming; however, determining whether a given triangulation is combinatorially regular seems to be very hard in general: a result by Richter-Gebert ([RG96] corollary 10.4.1) states that already in dimension 4 there are infinitely many minor-minimal non-combinatorially regular diagrams. In other words, combinatorial types of d-diagrams with n vertices may not be characterized by excluding a finite set of forbidden minors. We plan to investigate a weaker variant of this problem restricted to a subspace of combinatorially regular triangulations by focusing on rational embedding of diagrams that are easier to study for enumeration purposes. This natural restriction of the space of diagrams links our problem to the topic of tropical geometry, whose objects are rational polyhedral complexes in \mathbb{R}^d .

Acknowledgment: The authors wish to thank Fernando de Goes for his help on this project. Partial funding was generously provided by the National Science Foundation through grants CCF-1011944, CCF-0811373, and CMMI-0757106.

References

- [ACSYD] P. Alliez, D. Cohen-Steiner, M. Yvinec, and M. Desbrun. Variational tetrahedral meshing. In ACM SIGGRAPH 2005 Courses.
- [Aur87a] F. Aurenhammer. A criterion for the affine equivalence of cell complexes in \mathbb{R}^d and convex polyhedra in \mathbb{R}^{d+1} . Discrete and Computational Geometry, 2(1):49–64, 1987.
- [Aur87b] F. Aurenhammer. Recognising polytopical cell complexes and constructing projection polyhedra. Journal of Symbolic Computation, 3(3):249–255, 1987.
- [Ban90] Thomas F. Banchoff. Beyond the Third Dimension: Geometry, Computer Graphics, and Higher Dimensions. W. H. Freeman, 1990.
- [Bos98] Alain Bossavit. Computational Electromagnetism. Academic Press, Boston, 1998.
- [BP83] B. R. Baligaa and S. V. Patankarb. A Control Volume Finite-Element Method For Two-Dimensional Fluid Flow And Heat Transfer. Numerical Heat Transfer, 6:245–261, 1983.
- [CDL08] S.-W. Cheng, T. K. Dey, and J. Levine. Theory of a practical Delaunay meshing algorithm for a large class of domains. In B. Bhattacharya et al., editor, Algorithms, Architecture and Information Systems Security, volume 3 of World Scientific Review, pages 17–41, 2008.
- [CH80] R. Connelly and D.W. Henderson. A convex 3-complex not simplicially isomorphic to a strictly convex complex. In *Mathematical Proceed*ings of the Cambridge Philosophical Society, volume 88, pages 299–306. Cambridge Univ Press, 1980.
- [DKT07] Mathieu Desbrun, Eva Kanso, and Yiying Tong. Discrete differential forms for computational modeling. In A. Bobenko and P. Schröder, editors, *Discrete Differential Geometry*. Springer, 2007.
- [DLRS] JA De Loera, J. Rambau, and F. Santos. Triangulations: Applications, Structures and Algorithms. Algorithms and Computation in Mathematics, Springer Verlag, Heidelberg, to appear.
- [Ede87] H. Edelsbrunner. Algorithms in Combinatorial Geometry. Springer-Verlag, 1987.
- [ES86] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. Discrete Comput. Geom., 1:25–44, 1986.
- [Ewa96] G. Ewald. Combinatorial convexity and algebraic geometry. Springer Verlag, 1996.
- [FSBS06] Matthew Fisher, Boris Springborn, Alexander I. Bobenko, and Peter Schröder. An algorithm for the construction of intrinsic delaunay triangulations with applications to digital geometry processing. In ACM SIGGRAPH Courses, pages 69–74, 2006.
- [GKZ94] I.M. Gelfand, M.M. Kapranov, and A.V. Zelevinsky. Discriminants, resultants, and multidimensional determinants. Springer, 1994.

- [Gli05] D. Glickenstein. Geometric triangulations and discrete Laplacians on manifolds. Arxiv preprint math/0508188, 2005.
- [GP10] Leo J. Grady and Jonathan R. Polimeni. Discrete Calculus: Applied Analysis on Graphs for Computational Science. Springer, 2010.
- [Grü03] B. Grünbaum. Convex polytopes. Springer Verlag, 2003.
- [Lee91] C.W. Lee. Regular triangulations of convex polytopes. Applied Geometry and Discrete Mathematics—The Victor Klee Festschrift (P. Gritzmann, B. Sturmfels, eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Amer. Math. Soc, 4:443–456, 1991.
- [LWL⁺09] Y. Liu, W. Wang, B. Lévy, F. Sun, D.M. Yan, L. Lu, and C. Yang. On centroidal voronoi tessellationenergy smoothness and fast computation. ACM Transactions on Graphics (TOG), 28(4):1–17, 2009.
- [McC89] S. F. McCormick. Multilevel Adaptive Methods for Partial Differential Equations. SIAM, 1989.
- [Mjo06] Eric Mjolsness. The Growth and Development of Some Recent Plant Models: A Viewpoint. *Journal of Plant Growth Regulation*, 25(4):270–277, 2006.
- [MMdGD11] P. Mullen, P. Memari, F. de Goes, and M. Desbrun. HOT: Hodge-Optimized Triangulations. ACM Trans. Graph., 30(4):103:1–103:12, August 2011.
- [Mus97] O.R. Musin. Properties of the Delaunay triangulation. In Proceedings of the thirteenth annual symposium on Computational geometry, page 426. ACM, 1997.
- [OBSC00] Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. Spatial tessellations: Concepts and applications of Voronoi diagrams. Probability and Statistics. Wiley, 2nd edition, 2000.
- [PMH07] A. Paluszny, S. Matthäi, and M. Hohmeyer. Hybrid finite element finite volume discretization of complex geologic structures and a new simulation workflow demonstrated on fractured rocks. *Geofluids*, 7:186–208, 2007.
- [PS85] F. P. Preparata and M. I. Shamos. Computational Geometry: An Introduction. Springer-Verlag, 1985.
- [Raj94] VT Rajan. Optimality of the Delaunay triangulation in \mathbb{R}^d . Discrete and Computational Geometry, 12(1):189–202, 1994.
- [RG96] J. Richter-Gebert. Realization spaces of polytopes. Citeseer, 1996.
- [Sch03] A. Schwartz. Constructions of cubical polytopes. *PhD diss., TU Berlin*, 2003
- [Sha93] N.R. Shah. A parallel algorithm for constructing projection polyhedra. Information Processing Letters, 48(3):113–119, 1993.
- [Ste22] E. Steinitz. Polyeder und raumeinteilungen. Encyclopädie der mathematischen Wissenschaften, 3(9):1–139, 1922.
- [Tho06] R.R. Thomas. Lectures in geometric combinatorics. Amer Mathematical Society, 2006.
- [TWAD09] Jane Tournois, Camille Wormser, Pierre Alliez, and Mathieu Desbrun. Interleaving delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. ACM Trans. Graph., 28:75:1–75:9, July 2009.
- [Zie95] G.M. Ziegler. Lectures on polytopes. Springer, 1995.